

TraceTrans: Translation and Spatial Tracing for Surgical Prediction

Xiyu Luo^{1*}, Haodong Li^{1*}, Xinxing Cheng², He Zhao³, Yang Hu⁴, Xuan Song⁵, Tianyang Zhang^{6†}

¹Southern University of Science and Technology, China

²University of Birmingham, United Kingdom

³University of Liverpool, United Kingdom

⁴University of Leicester, United Kingdom

⁵Jilin University, China

⁶University of Oxford, United Kingdom

luoxy2022@mail.sustech.edu.cn, tianyang.zhang@eng.ox.ac.uk

Abstract

Image-to-image translation models have achieved notable success in converting images across visual domains and are increasingly used for medical tasks such as predicting post-operative outcomes and modeling disease progression. However, most existing methods primarily aim to match the target distribution and often neglect spatial correspondences between the source and translated images. This limitation can lead to structural inconsistencies and hallucinations, undermining the reliability and interpretability of the predictions. These challenges are accentuated in clinical applications by the stringent requirement for anatomical accuracy. In this work, we present TraceTrans, a novel deformable image translation model designed for post-operative prediction that generates images aligned with the target distribution while explicitly revealing spatial correspondences with the pre-operative input. The framework employs an encoder for feature extraction and dual decoders for predicting spatial deformations and synthesizing the translated image. The predicted deformation field imposes spatial constraints on the generated output, ensuring anatomical consistency with the source. Extensive experiments on medical cosmetology and brain MRI datasets demonstrate that TraceTrans delivers accurate and interpretable post-operative predictions, highlighting its potential for reliable clinical deployment.

Code — <https://github.com/xyluoxy/TraceTrans>

Introduction

Medical image-to-image translation is particularly important for post-operative outcome prediction, where the objective is to generate realistic post-operative images from pre-operative inputs while preserving anatomical plausibility. Such predictions help clinicians visualize potential surgical outcomes, plan interventions, and communicate expected changes to patients. Beyond cosmetic surgery, similar approaches can model longitudinal changes in neurological diseases or tumor evolution, offering valuable insights for both clinical practice and research.

*These authors contributed equally.

†Corresponding author.

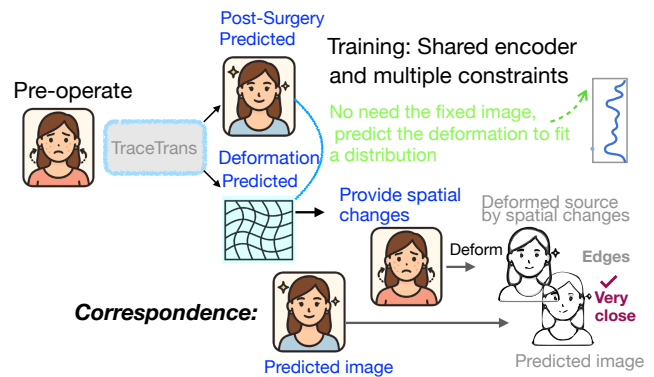


Figure 1: TraceTrans employs an encoder with dual decoders to predict both spatial deformation and translated images. The deformation field provides pixel-level correspondences from the source to the target, ensuring structural alignment while adapting to the target distribution. No fixed reference image is required during training and inference.

Recent advances in deep generative models, including GAN-based (Isola et al. 2016; Zhu et al. 2017) and diffusion-based approaches (Saharia et al. 2021; Rombach et al. 2021; Su et al. 2023), have improved image realism across translation tasks. However, most existing methods focus on matching the target distribution and neglect spatial correspondences between source and translated images, leading to inconsistencies and hallucinations that compromise reliability and interpretability. A common workaround is to enforce structural invariance between input and output to improve consistency (Zhang et al. 2019), but this conflicts with the objective of post-operative prediction, where structural changes must be realistically modeled.

To address the challenges of structural inconsistency and provide traceable spatial changes during translation for post-operative prediction, we aim to develop a conditional generative framework capable of jointly predicting spatial deformations and post-operative translated outcomes, as illustrated in Figure 1. Conventional registration approaches, such as VoxelMorph (Balakrishnan et al. 2018), are effective at estimating deformation fields between a fixed and a mov-

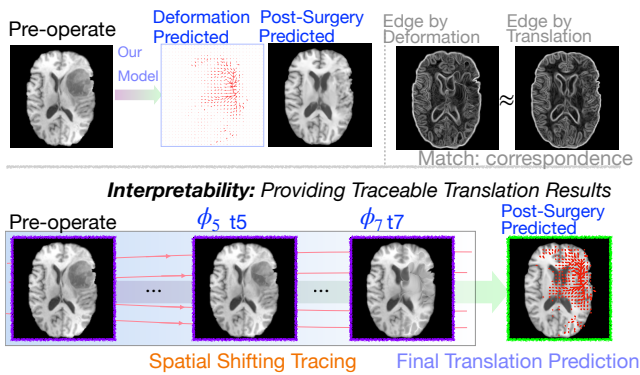


Figure 2: TraceTrans jointly predicts spatial deformations and post-operative translations, enabling pixel-level correspondences and traceable structural changes for interpretable surgical outcome prediction. The second row illustrates the translation process, where spatial shifts are composed from velocity fields, while the first row shows that the close alignment between edges of the deformed pre-operative based on the predicted deformation and predicted post-operative images confirms anatomical correspondence.

ing image, yet they require a fixed reference image, which is unavailable in translation tasks where only the source image is provided. Moreover, these methods focus solely on deformation estimation and cannot generate images that align with the target distribution. In contrast, our goal integrates deformation prediction and image synthesis within a unified end-to-end network. This design allows the model to generate anatomically consistent translations while simultaneously providing pixel-level correspondences between the translated and source images, thereby enhancing both interpretability and reliability in clinical applications such as surgical outcome prediction.

Therefore, we propose TraceTrans, a novel two-stream end-to-end model for surgical outcome prediction. To the best of our knowledge, this is the first approach to generate predicted post-operative images with spatial traceability without requiring fixed reference images as additional inputs. The model consists of an encoder for feature extraction and two decoders that predict a velocity field and synthesize the translated image. The velocity field is integrated into a deformation field through an iterative composition process, and the resulting spatial changes are illustrated in Figure 2. To preserve anatomical structure in the translated images, multiple structural constraints are incorporated into the training objective, as detailed in Sec. Method. We evaluate TraceTrans on both medical cosmetology and BraTS-Reg (Baheti et al. 2021) datasets, which contain pre-operative–post-operative face image pairs and brain MRI slice pairs. Quantitative and qualitative comparisons with existing image translation models demonstrate that TraceTrans produces higher-quality translations with superior structural consistency to the source images, as presented in Sec. Experiments.

Our main contributions can be summarized as follows.

- We define the interpretable surgical prediction task as generating spatially traceable post-operative images solely from pre-operative inputs.
- We propose TraceTrans, a novel two-stream end-to-end model for surgical outcome prediction. To the best of our knowledge, this is the first approach to generate predicted post-operative images with spatial traceability, without requiring fixed reference images during translation training to capture structural changes.
- TraceTrans is evaluated on two representative scenarios: predicting facial structural changes after cosmetic surgery and modeling longitudinal brain MRI changes in glioma patients, with results demonstrating superior effectiveness compared to prior approaches.

Related Works

Image-to-Image Translation

Image-to-image translation is a core task in computer vision and aims to generate image in target domain based on image in source domain while preserving the structure of subject in the image. From the perspective of dataset, current image-to-image translation methods can be roughly categorized into supervised methods and unsupervised methods that require datasets with paired data and two datasets in different domains, respectively, and from the perspective of architecture, Generative Adversarial Network (GAN) (Goodfellow et al. 2014) and diffusion model (Ho, Jain, and Abbeel 2020) are the most popular basic architectures for this task, and many works are adapted from one of them.

Pix2Pix (Isola et al. 2016) and CycleGAN (Zhu et al. 2017) are pioneering works in supervised GAN and unsupervised GAN and both lay a solid foundation for the development of GAN-based image-to-image translation models (Wang et al. 2017; Fu et al. 2018; Kong et al. 2021; Amirkolae and Amirkolae 2022; Xin et al. 2024). Palette (Saharia et al. 2021) is the first systematic and diffusion-based image translation framework that achieves outstanding performance in multiple image translation tasks. After the success of Palette, many diffusion-based image-to-image translation models are proposed, such as Latent Diffusion Models (Rombach et al. 2021), Diffi2I (Xia et al. 2023), DDIB (Su et al. 2023) and so on. Although these models are capable of generating good translated images, it is difficult to maintain a precise structure correspondence between the translated image and the source image. Hence, they cannot be directly applied to surgical prediction task.

Surgical Prediction

Prediction of the natural progression of disease and prediction of surgical outcome are two important tasks that facilitate the diagnosis and planning of surgery. With the advancement of image-to-image translation models, many works adapt them, such as GAN-based models (Ravi et al. 2022; Xia et al. 2021; Yoo, Choi, and Kim 2020; Hwang et al. 2021) and diffusion-based models (Yoon et al. 2023) to directly synthesize the outcome image and achieve excellent results in their specific field. We choose the GAN model as our backbone due to its high response ability, and our aim

is to provide a more robust and reliable surgical outcome prediction model.

Medical Image Registration

Medical image registration aims to establish the spatial correspondence between two images that shares the same structure for the subject and is the crucial and fundamental step in medical image analysis. Traditional methods solve the best registration by optimizing a predefined similarity metric (Rueckert et al. 1999; Avants et al. 2011) and the optimization process is relatively slow. Deep learning-based methods, such as VoxelMorph (Balakrishnan et al. 2018), accelerate the prediction process by modeling registration as a regression problem. Due to the outstanding performance of VoxelMorph, many follow-up works (Kim et al. 2020; Mok and Chung 2020; Cheng et al. 2025) are adapted from it. Although these methods are capable of predicting the registration with high precision, they cannot be applied directly to surgical prediction, since these models require fixed images as input, and the fixed images are unavailable before prediction.

Method

The backbone of TraceTrans adopts a dual-stream GAN architecture, consisting primarily of an encoder ϵ , followed by a pixel translation decoder g , and a parallel deformation flow decoder f . Two corresponding discriminators are employed to distinguish between the translated and deformed outputs, respectively. An overview of TraceTrans’s architecture is illustrated in Figure 3. We begin by defining the deformation process in Sec. Image Deformation, then describe the structure of the model’s backbone in Sec. Model Architecture, and finally provide details of the training procedure in Sec. Training.

Image Deformation

Given m as the n -dimensional source (moving) image defined over grid $\Omega \subset \mathbb{R}^n$, let $\mathbf{u} : \Omega \rightarrow \mathbb{R}^n$ and $\mathbf{v} : \Omega \rightarrow \mathbb{R}^n$ be the displacement field and the velocity field, respectively. To encourage a diffeomorphic deformation, we integrate the predicted stationary velocity field \mathbf{v} to form the displacement field \mathbf{u} using scaling and squaring layers (Ashburner 2007). The displacement field \mathbf{u}_T at any time T is integrated from the velocity field \mathbf{v} according to

$$\mathbf{u}_T(\mathbf{p}) = \int_0^T \mathbf{v}((Id + \mathbf{u}_t)(\mathbf{p})) dt. \quad (1)$$

In our code implementation, we adopt $T = 1$ to obtain the final displacement field $\mathbf{u} = \mathbf{u}_1$, using a 7-step integration. Consequently, the resulting deformation field ϕ can be expressed as

$$\phi(\mathbf{p}) = (Id + \mathbf{u})(\mathbf{p}), \quad (2)$$

where $\mathbf{p} \in \Omega$ is an arbitrary spatial point of dimension n in the moving image m , which is, in the 2D case, a pixel. $Id(\mathbf{p}) = \mathbf{p}$ is an identity deformation. Once the deformation field ϕ is predicted, the warped image m' can be generated by $m' = m \circ \phi$ if m is defined over \mathbb{R}^n .

Note that $m \circ \phi(\mathbf{p})$ is not defined for $\phi(\mathbf{p}) \notin \Omega$, interpolation is needed to estimate the value of m at $\phi(\mathbf{p}) \notin \Omega$. Hence the spatial transformer network (Jaderberg et al. 2015) is utilized to perform a differentiable linear interpolation at a point $\phi(\mathbf{p}) \notin \Omega$. More specifically, the linear interpolation process is defined as

$$m'(\mathbf{p}) \approx \sum_{\mathbf{q} \in Z(\mathbf{p}')} m(\mathbf{q}) \prod_{i=1}^n (1 - |\mathbf{q}_i - \mathbf{p}'_i|), \quad (3)$$

where m' is the warped image, $\mathbf{p}' = \phi(\mathbf{p})$, and $Z(\mathbf{p}')$ is the set of 2^n discrete neighbor points of \mathbf{p}' .

Model Architecture

As illustrated in Figure 3, the network comprises three components: 1) an encoder ϵ hierarchically extracts multi-scale features from the pre-operative images; 2) two decoders, g and f , are employed to predict pixel translations of post-operative images and to estimate spatial transformations, respectively; 3) two discriminators, D_{trans} and D_{warp} , are used to distinguish the results of translated post-operative and the deformed pre-operative images, respectively.

Given a pre-operative image \mathbf{x} , the encoder ϵ extracts multi-scale features $\epsilon(\mathbf{x})$. Then, the translation decoder g takes these multi-scale features to generate the translated post-operative image as:

$$\mathbf{y}_{\text{trans}} = g(\epsilon(\mathbf{x})). \quad (4)$$

In the Spatial Transformation Module, we note that the warped result \mathbf{y}_{warp} is not directly generated by the deformation decoder f . Instead, f produces a velocity field \mathbf{v} :

$$\mathbf{v} = f(\epsilon(\mathbf{x})). \quad (5)$$

The final deformation field ϕ is then obtained by integrating the velocity field \mathbf{v} , as defined in Eq. 1 and Eq. 2. Next, ϕ is applied to warp the pre-operative image \mathbf{x} , producing the deformed result:

$$\mathbf{y}_{\text{warp}} = \mathbf{x} \circ \phi. \quad (6)$$

Suppose the ground truth post-operative reference image is \mathbf{y} . The translated $\mathbf{y}_{\text{trans}}$ is combined with \mathbf{y} to form a fake-real pair, and fed to the translation discriminator D_{trans} . Similarly, \mathbf{y}_{warp} is also combined with \mathbf{y} and fed to the warp discriminator D_{warp} . Our discriminators adopt the conditional GAN formulation, in which the source image \mathbf{x} is used as a condition to guide the discrimination process (Mirza and Osindero 2014) and are discussed in the next Section.

To ensure structural correspondence between $\mathbf{y}_{\text{trans}}$ and \mathbf{y}_{warp} , a cross-domain constraint is applied between the two images. This constraint is implemented via differentiable normalized mutual information (DNMI) loss.

Training

DNMI We adopt differentiable mutual information to compute the loss between generated \mathbf{y}_{warp} and reference \mathbf{y} , as well as the cross-domain similarity between $\mathbf{y}_{\text{trans}}$ and \mathbf{y}_{warp} , in order to include the cross-domain constraint in the

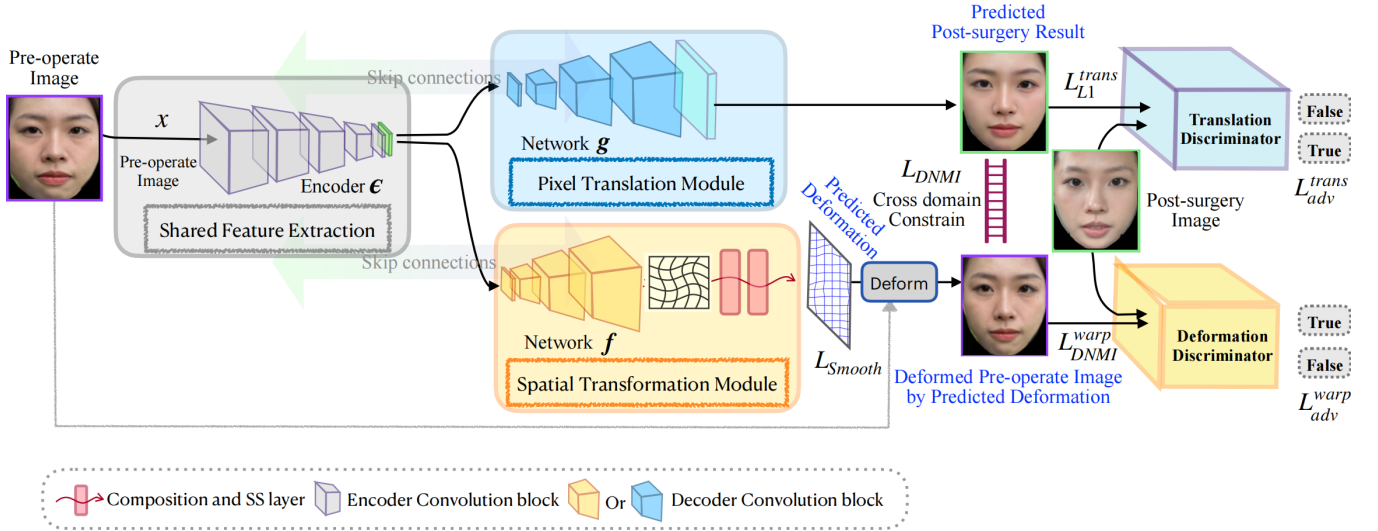


Figure 3: Overview of the TraceTrans framework. An encoder extracts features from the pre-operative image, which are processed by two decoders: the pixel translation module generates the post-operative prediction, and the spatial transformation module predicts the deformation field. The deformation ensures structural correspondence between the pre-operative and post-operative domains, while discriminators and multi-level constraints jointly optimize translation fidelity and anatomical consistency.

final backpropagation. The definition of normalized differentiable mutual information (Qiu et al. 2021) is given as follows.

$$NMI(X, Y) = \frac{H(X) + H(Y)}{H(X, Y)} \quad (7)$$

where X, Y are two random variables, $H(X)$ is the entropy. Note that NMI is not differentiable, Parzen window (Thevenaz and Unser 2000) is used to substitute for original rectangular window. Since the more two images are structurally similar to each other, the higher the DNMI is, $\mathcal{L}_{DNMI}(x, y)$ is formulated as $\mathcal{L}_{DNMI}(x, y) = -DNMI(x, y)$.

Trans-Deform Ratio In order to clearly define the computation of losses of our model, it is necessary to explain the controllable ratio of the two GAN streams in TraceTrans. Since the two decoders g and f are responsible for translation and deformation tasks respectively, we can control the backpropagation proportions of the two tasks. We define the proportion of the translation task in TraceTrans as α , and the deformation task $1 - \alpha$. α is a hyperparameter that can be adjusted during training.

A content alignment loss \mathcal{L}_{align} is defined to measure the semantic consistency between the generated outputs and their corresponding targets, as well as the cross-domain constraint:

$$\mathcal{L}_{align} = \alpha \mathcal{L}_{LI}^{trans} + (1 - \alpha) \mathcal{L}_{DNMI}^{warp} + \min(\alpha, 1 - \alpha) \cdot \gamma \cdot \mathcal{L}_{DNMI}^{cross} \quad (8)$$

where $\mathcal{L}_{LI}^{trans} = \|\mathbf{y}_{trans} - \mathbf{y}\|_1$, $\mathcal{L}_{DNMI}^{warp} = -DNMI(\mathbf{y}_{warp}, \mathbf{y})$, and $\mathcal{L}_{DNMI}^{cross} = -DNMI(\mathbf{y}_{trans}, \mathbf{y}_{warp})$. Here, γ is used to control the strength of cross-domain constraint. The sensitivity analysis of γ is discussed in Sec: Cross-Domain Constraint.

Training Framework The adversarial loss $\mathcal{L}_{adv, G}$ for the generator G and $\mathcal{L}_{adv, D}$ for the discriminator D are defined by the minimax game,

$$\theta_G^* = \arg \min_{\theta_G} \max_{\theta_D} \left[\mathbb{E}_{x \in X} |D(x, y(x))|^2 + \mathbb{E}_{x \in X} |1 - D(x, G(x))|^2 \right], \quad (9)$$

and written as follows, respectively.

$$\mathcal{L}_{adv, G}(\theta_G) = |D(x, G(x)) - 1|^2 \quad (10)$$

$$\mathcal{L}_{adv, D}(\theta_D) = \frac{1}{2} (|D(x, y(x)) - 1|^2 + |D(x, G(x))|^2) \quad (11)$$

where θ_G, θ_D denote the parameters for G, D respectively; θ_G^* denotes the best parameter for G ; X is the set of source images, $y(x)$ is the target image corresponding to the source image x .

Overall Loss Function Given the trans-proportion α and the content alignment loss \mathcal{L}_{align} , the overall generator loss \mathcal{L}_G is computed as

$$\mathcal{L}_G = \mathcal{L}_{align} + \lambda_{adv} (\alpha \mathcal{L}_{adv, G}^{trans} + (1 - \alpha) \mathcal{L}_{adv, G}^{warp}) + (1 - \alpha) \lambda_{smooth} \mathcal{L}_{smooth}, \quad (12)$$

and the overall discriminator loss \mathcal{L}_D is

$$\mathcal{L}_D = \lambda_{adv} (\alpha \mathcal{L}_{adv, D}^{trans} + (1 - \alpha) \mathcal{L}_{adv, D}^{warp}). \quad (13)$$

The $\mathcal{L}_{adv, G}^{trans}$ and $\mathcal{L}_{adv, G}^{warp}$ represent the adversarial losses for the two generation tasks, respectively. Likewise $\mathcal{L}_{adv, D}^{trans}$ and $\mathcal{L}_{adv, D}^{warp}$ are adversarial losses for two discriminators. The smoothness term $\mathcal{L}_{smooth}(\mathbf{v}) = \sum_{\mathbf{p} \in \Omega} \|\nabla \mathbf{v}(\mathbf{p})\|^2$ is used to regularize the spatial gradients of the deformation field and to ensure the diffeomorphism of the velocity field.

Method	SSIM(%) \uparrow	MAE \downarrow	PSNR \uparrow	NMI \uparrow
MUNIT (Huang et al. 2018)	79.64 (± 4.11)	22.65 (± 10.57)	28.96 (± 2.90)	1.1670 (± 0.0199)
CycleGAN (Zhu et al. 2017)	78.51 (± 4.63)	17.92 (± 4.50)	29.33 (± 1.93)	1.1802 (± 0.0333)
GcGAN (Fu et al. 2018)	77.67 (± 4.55)	18.81 (± 5.21)	29.26 (± 1.97)	1.1738 (± 0.0314)
Pix2Pix (Isola et al. 2016)	79.93 (± 4.20)	16.72 (± 3.98)	29.60 (± 1.08)	1.1821 (± 0.0446)
Palette (Saharia et al. 2021)	69.80 (± 4.55)	25.38 (± 7.55)	27.63 (± 2.25)	1.1737 (± 0.0242)
DDIB (Su et al. 2023)	74.63 (± 5.72)	20.47 (± 8.53)	28.32 (± 2.44)	1.1745 (± 0.0243)
TraceTrans	82.92 (± 4.14)	14.37 (± 4.41)	29.86 (± 0.93)	1.1943 (± 0.0414)

Table 1: Quantitative comparison on the Face Cosmetology datasets. TraceTrans obtained the best result in all the metrics on this datasets.

Method	SSIM(%) \uparrow	MAE \downarrow	PSNR \uparrow	NMI \uparrow
MUNIT (Huang et al. 2018)	72.10 (± 7.35)	13.35 (± 4.31)	31.52 (± 2.39)	1.2144 (± 0.0235)
CycleGAN (Zhu et al. 2017)	72.10 (± 7.59)	14.95 (± 4.59)	31.45 (± 2.58)	1.2080 (± 0.0218)
GcGAN (Fu et al. 2018)	72.67 (± 8.24)	15.06 (± 4.92)	31.28 (± 2.72)	1.2100 (± 0.0240)
Pix2Pix (Isola et al. 2016)	74.01 (± 6.85)	14.74 (± 3.91)	31.30 (± 0.48)	1.2112 (± 0.0291)
Palette (Saharia et al. 2021)	61.96 (± 10.23)	21.19 (± 8.03)	29.10 (± 2.48)	1.2055 (± 0.0262)
DDIB (Su et al. 2023)	65.47 (± 6.46)	18.31 (± 4.26)	31.18 (± 1.74)	1.1864 (± 0.0146)
TraceTrans	75.97 (± 6.90)	13.30 (± 3.93)	31.48 (± 0.44)	1.2181 (± 0.0285)

Table 2: Quantitative comparison on the Brain MRI datasets. TraceTrans obtained the best result in the majority of metrics on this datasets.




Figure 4: Qualitative comparison of our method and several baselines. The images in the first row and the second row are the results on face cosmetology dataset and brain MRI dataset, respectively. The first two columns show the input and ground truth images, the last column shows Sobel edge maps derived from the deformed source images using the predicted deformation field of our method.

Experiments

We conduct several experiments and an ablation study to evaluate the effectiveness of our method in surgical prediction, and a sensitivity study to find the best value for hyper-parameters α and γ . Two different domains are considered: face cosmetic surgery and brain tumor surgery. We compare our model with several representative models including both GAN-based and diffusion-based models.

Datasets

Face Cosmetology This dataset is a proprietary collection curated by us and consists of 412 paired frontal face images captured before and after cosmetic surgery. Each pair corresponds to the same individual and was collected with informed consent from the image providers. To ensure consistency

across samples, an automated face detection pipeline is applied to locate facial regions in each image, followed by masking of non-facial areas such as background and hair. The cropped facial region is then resized to a standardized resolution of $256 \times 256 \times 3$, where the three channels correspond to RGB color information. After preprocessing, 330 image pairs are used for training and the remaining 82 pairs are reserved for testing.

Brain MRI We also evaluate our method on the BraTS-Reg (Baheti et al. 2021) dataset, which provides pre-operative and follow-up MRI scans for patients with diffuse glioma. The dataset includes multimodal MRI sequences (T1, T2, FLAIR, and T1 contrast-enhanced) collected longitudinally for each subject. To mitigate inter-scan variability caused by different acquisition protocols, histogram stan-

standardization (Nyul, Udupa, and Zhang 2000) is first applied to the pre-operative and follow-up T1 and T2 scans. We then perform Z-normalization, rescale the intensity to the $[0, 1]$ range, and crop the scans to a resolution of $192 \times 192 \times 128$. The pre-operative and follow-up scans are rigidly aligned using the SimpleITK registration framework (Yaniv et al. 2018). For experiments, 160 subjects in BraTS-Reg are split into training and testing sets with a 7:3 ratio and the central eleven slices are extracted to form a dataset with 1232 training pairs and 528 testing pairs.

Implementation Details

Our model is implemented in PyTorch and trained for 500 epochs using a batch size of 8. All experiments are conducted on an NVIDIA A100 GPU with 40GB of VRAM. Input images are resized to 256×256 and normalized to the range $[-1, 1]$. The network is optimized using the Adam optimizer with a learning rate of 2×10^{-4} . The encoder contains 5 downsampling layers, where the number of output channels for each layer is $[64, 128, 256, 512, 512]$. We use the PatchGAN discriminators in our network as proposed in Pix2Pix (Isola et al. 2016). The optimal hyperparameters used for training are $\alpha = 0.5$, $\gamma = 1.0$, $\lambda_{adv} = 0.01$, and $\lambda_{smooth} = 0.2$. We use 16 bins for $\mathcal{L}_{DNMI}^{warp}$ and 32 bins for $\mathcal{L}_{DNMI}^{cross}$.

Evaluation Metrics

We evaluate the performance of TraceTrans and baseline methods using four standard image similarity metrics: Structural Similarity Index (SSIM) (Wang et al. 2004), Peak Signal-to-Noise Ratio (PSNR), Mean Absolute Error (MAE), and Normalized Mutual Information (NMI) (Studholme, Hill, and Hawkes 1999). Higher values of SSIM, PSNR, and NMI, and a lower value of MAE, indicate better image quality and closer alignment with the ground truth.

Experiment Results on Face Cosmetology

Performance evaluations on the face cosmetology dataset are presented in Table 1, where TraceTrans outperforms all other models across all metrics. The significant improvements in SSIM and NMI indicate TraceTrans’s superior ability to preserve structural correspondence. In addition, the increase in PSNR and reduction in MAE suggest that the predicted post-operative images generated by TraceTrans exhibit better texture translation quality compared to other models. The qualitative results are shown in the first row of Figure 4, where the predicted post-operative image produced by TraceTrans is visually more similar to the reference image than those generated by the other methods, aligning well with the quantitative results.

Experiment Results on Brain MRI

To further demonstrate the effectiveness of TraceTrans, we additionally evaluated the model on Brain MRI dataset. The quantitative results are shown in Table 2 and TraceTrans performs best in SSIM, MAE, NMI and slightly worse than MUNIT in PSNR. Since the difference between TraceTrans

and MUNIT on PSNR is less than 5% of either standard deviation, then this difference is negligible. TraceTrans shows significant SSIM improvement while matching MUNIT’s PSNR and achieving the lowest MAE, demonstrating balanced excellence in both structural and textural translation. Qualitative results are illustrated in the second row of Figure 4, the image generated by our model is visually best similar to the reference image.

Ablation Study

To study the necessity of each component, we conducted two ablation studies: one evaluating structural correspondence and the other assessing fidelity of the predicted results, with further details provided in the sensitivity analysis.




Figure 5: Illustration of correspondence comparison in the ablation study. The proposed two-stream constrained network generates deformations that align with the translated images, whereas a similar single-stream architecture fails to preserve this correspondence.

Correspondence (evaluated by edge)	SSIM (%) \uparrow	PSNR \uparrow	NMI \uparrow
Single-stream (two networks)	55.58	18.24	0.2167
Two-stream no constraints	58.50	19.21	0.2224
TraceTrans	59.48	19.34	0.2249

Table 3: Quantitative comparison of correspondence performance for different model components.

The evaluation of structural correspondence, which is the most critical aspect, is performed using deformations predicted solely from preoperative images. Specifically, the source image is warped using the predicted deformation field, and Sobel edge maps are computed from the warped image. These are then compared with the Sobel edge maps of the predicted postoperative images to quantify the difference between corresponding edges. For this ablation, three configurations are examined: (1) a single-stream network that generates postoperative images and deformations in parallel without a shared encoder; (2) a two-stream network with a shared encoder but without the structural correspondence constraint; and (3) the proposed full model.

Table 3 compares correspondence performance across model variants. The single-stream design without shared encoding or correspondence constraints yields the lowest scores, while adding a two-stream structure improves all




Figure 6: Brain MRI predictions under $\alpha = 1, 0.75, 0.5$ and 0.25 .

metrics. The full TraceTrans model, which also includes the correspondence constraint, achieves the best results, confirming its effectiveness in preserving edge-level alignment between preoperative and postoperative predictions.

Additionally, an ablation on image fidelity (Table 4) shows that adding the deformation stream improves MAE, and further introducing the cross-domain constraint achieves the best fidelity, confirming the benefit of the full TraceTrans design.

Trans Stream	Deform Stream	Cross-Domain Constraint	Face MAE↓
✓	✗	✗	12.88
✓	✓	✗	12.03
✓	✓	✓	10.64

Table 4: Ablation study on image fidelity.

Sensitivity Analysis

Trans-Deform Ratio α controls the contribution of the two streams in our model. We evaluate its impact on two datasets, focusing on (1) traceability from input to predicted output and (2) the quality of the prediction.

To assess traceability, we use ANTs (Avants et al. 2011) to compute a deformation field from the pre-operative to the predicted post-operative image. We then measure similarity between the predicted image and the warped pre-operative image. Since the deformation stream is disabled when $\alpha = 1$, its field is excluded for fairness. At $\alpha = 1$, the model reduces to Pix2Pix, which often produces blurry MRI outputs, making ANTs registration unstable—though it remains stable on the face cosmetology dataset.

Although both models perform similarly in global facial regions, Pix2Pix struggles to reconstruct key features (e.g., eyes, nose, mouth). We address this by masking non-essential areas and computing edge Dice scores within the masked regions. Results are reported in the first row of Table 5. These findings show that the deformation stream enhances both spatial traceability and prediction quality.

We assess output quality by computing the SSIM between predicted and ground-truth post-operative brain MRIs. As shown in Table 5, $\alpha = 0.5$ provides the best trade-off between pixel-level similarity and structural consistency. Figure 6 presents qualitative comparisons across different α values.

Cross-Domain Constraint γ is a key parameter that controls the strength of the cross-domain constraint between

α	1	0.75	0.5	0.25
Masked Dice(%) (face)	44.72	49.26	51.10	49.62
SSIM(%) (MRI)	74.01	75.11	75.97	75.57

Table 5: Effect of different (α) values on traceability and prediction quality. Traceability is measured by Masked Dice between warped input (using ANTs) and predicted output.

y_{trans} and y_{warp} . Although both decoders share the same encoder, it is important to assess whether this constraint is necessary to ensure structural consistency between their outputs. We conduct a sensitivity analysis on five γ values ranging from 0 to 1 using the face cosmetology dataset.

The cross-domain constraint primarily enforces structural alignment between y_{trans} and y_{warp} , which we evaluate using SSIM. Within the tested range, larger γ values consistently improve structural consistency. Additionally, we observe that stronger constraints enhance traceability from the input to the predicted output, as reflected by higher masked Dice scores. Results are summarized in Table 6.

γ	0	0.25	0.5	0.75	1.0
Masked SSIM(%)	83.00	83.19	83.69	85.03	85.16
Masked Dice(%)	45.60	45.70	46.79	48.96	51.10

Table 6: Effect of Cross-Domain Constraint Strength (γ) on structural correspondence and traceability. Structural correspondence between y_{trans} and y_{warp} is measured by SSIM between unmasked facial areas of these two. Masked Dice is calculated in the same way as Table 5.

Conclusion

In this work, we introduced TraceTrans, a deformable image-to-image translation framework for postoperative prediction. Unlike prior methods that focus on distribution alignment and neglect spatial correspondence, TraceTrans predicts deformation fields to trace structural changes and ensure anatomical consistency. By integrating deformation prediction and image synthesis in an end-to-end design, it produces translations that are both visually accurate and interpretable. Experiments on medical cosmetology and brain MRI datasets show superior translation quality and structural fidelity, demonstrating its potential for clinical applications requiring precise postoperative modeling.

References

- Amirkolae, H. A.; and Amirkolae, H. A. 2022. Medical image translation using an edge-guided generative adversarial network with global-to-local feature fusion. *Journal of Biomedical Research*, 36(6): 409–422.
- Ashburner, J. 2007. A fast diffeomorphic image registration algorithm. *Neuroimage*, 38(1): 95–113.
- Avants, B. B.; Tustison, N. J.; Song, G.; Cook, P. A.; Klein, A.; and Gee, J. C. 2011. A reproducible evaluation of ANTs similarity metric performance in brain image registration. *NeuroImage*, 54(3): 2033–2044.
- Baheti, B.; Waldmannstetter, D.; Chakrabarty, S.; Akbari, H.; Bilello, M.; Wiestler, B.; Schwarting, J.; Calabrese, E.; Rudie, J.; Abidi, S.; Mousa, M.; Villanueva-Meyer, J.; Marcus, D. S.; Davatzikos, C.; Sotiras, A.; Menze, B.; and Bakas, S. 2021. The Brain Tumor Sequence Registration Challenge: Establishing Correspondence between Pre-Operative and Follow-up MRI scans of diffuse glioma patients. arXiv:2112.06979.
- Balakrishnan, G.; Zhao, A.; Sabuncu, M. R.; Guttag, J. V.; and Dalca, A. V. 2018. VoxelMorph: A Learning Framework for Deformable Medical Image Registration. *CoRR*, abs/1809.05231.
- Cheng, X.; Zhang, T.; Lu, W.; Meng, Q.; Frangi, A. F.; and Duan, J. 2025. SACB-Net: Spatial-awareness Convolutions for Medical Image Registration. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 5227–5237.
- Fu, H.; Gong, M.; Wang, C.; Batmanghelich, K.; Zhang, K.; and Tao, D. 2018. Geometry-Consistent Adversarial Networks for One-Sided Unsupervised Domain Mapping. *CoRR*, abs/1809.05852.
- Goodfellow, I. J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative Adversarial Networks. arXiv:1406.2661.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. *CoRR*, abs/2006.11239.
- Huang, X.; Liu, M.-Y.; Belongie, S.; and Kautz, J. 2018. Multimodal Unsupervised Image-to-Image Translation. arXiv:1804.04732.
- Hwang, D.-Y.; Choi, S.-H.; Shin, J.; Kim, M.; and Choi, Y.-H. 2021. GAN-Based ROI Image Translation Method for Predicting Image after Hair Transplant Surgery. *Electronics*, 10(24).
- Isola, P.; Zhu, J.; Zhou, T.; and Efros, A. A. 2016. Image-to-Image Translation with Conditional Adversarial Networks. *CoRR*, abs/1611.07004.
- Jaderberg, M.; Simonyan, K.; Zisserman, A.; and Kavukcuoglu, K. 2015. Spatial Transformer Networks. *CoRR*, abs/1506.02025.
- Kim, B.; Kim, D. H.; Park, S. H.; Kim, J.; Lee, J.-G.; and Ye, J. C. 2020. CycleMorph: Cycle Consistent Unsupervised Deformable Image Registration. arXiv:2008.05772.
- Kong, L.; Lian, C.; Huang, D.; Li, Z.; Hu, Y.; and Zhou, Q. 2021. Breaking the Dilemma of Medical Image-to-image Translation. arXiv:2110.06465.
- Mirza, M.; and Osindero, S. 2014. Conditional Generative Adversarial Nets. arXiv:1411.1784.
- Mok, T. C. W.; and Chung, A. C. S. 2020. Fast Symmetric Diffeomorphic Image Registration with Convolutional Neural Networks. *CoRR*, abs/2003.09514.
- Nyul, L.; Udupa, J.; and Zhang, X. 2000. New variants of a method of MRI scale standardization. *IEEE Transactions on Medical Imaging*, 19(2): 143–150.
- Qiu, H.; Qin, C.; Schuh, A.; Hammernik, K.; and Rueckert, D. 2021. Learning Diffeomorphic and Modality-invariant Registration using B-splines. In Heinrich, M.; Dou, Q.; de Bruijne, M.; Lellmann, J.; Schläfer, A.; and Ernst, F., eds., *Proceedings of the Fourth Conference on Medical Imaging with Deep Learning*, volume 143 of *Proceedings of Machine Learning Research*, 645–664. PMLR.
- Ravi, D.; Blumberg, S. B.; Ingala, S.; Barkhof, F.; Alexander, D. C.; and Oxtoby, N. P. 2022. Degenerative adversarial neuroimage nets for brain scan simulations: Application in ageing and dementia. *Medical Image Analysis*, 75: 102257.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2021. High-Resolution Image Synthesis with Latent Diffusion Models. *CoRR*, abs/2112.10752.
- Rueckert, D.; Sonoda, L.; Hayes, C.; Hill, D.; Leach, M.; and Hawkes, D. 1999. Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Transactions on Medical Imaging*, 18(8): 712–721.
- Saharia, C.; Chan, W.; Chang, H.; Lee, C. A.; Ho, J.; Salimans, T.; Fleet, D. J.; and Norouzi, M. 2021. Palette: Image-to-Image Diffusion Models. *CoRR*, abs/2111.05826.
- Studholme, C.; Hill, D.; and Hawkes, D. 1999. An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognition*, 32(1): 71–86.
- Su, X.; Song, J.; Meng, C.; and Ermon, S. 2023. Dual Diffusion Implicit Bridges for Image-to-Image Translation. arXiv:2203.08382.
- Thevenaz, P.; and Unser, M. 2000. Optimization of mutual information for multiresolution image registration. *IEEE Transactions on Image Processing*, 9(12): 2083–2099.
- Wang, T.; Liu, M.; Zhu, J.; Tao, A.; Kautz, J.; and Catanzaro, B. 2017. High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs. *CoRR*, abs/1711.11585.
- Wang, Z.; Bovik, A.; Sheikh, H.; and Simoncelli, E. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612.
- Xia, B.; Zhang, Y.; Wang, S.; Wang, Y.; Wu, X.; Tian, Y.; Yang, W.; Timotfe, R.; and Gool, L. V. 2023. DiffI2I: Efficient Diffusion Model for Image-to-Image Translation. arXiv:2308.13767.
- Xia, T.; Chartsias, A.; Wang, C.; and Tsiftaris, S. A. 2021. Learning to synthesise the ageing brain without longitudinal data. *Medical Image Analysis*, 73: 102169.
- Xin, B.; Young, T.; Wainwright, C. E.; Blake, T.; Lebrat, L.; Gaass, T.; Benkert, T.; Stemmer, A.; Coman, D.; and

- Dowling, J. 2024. Deformation-aware GAN for Medical Image Synthesis with Substantially Misaligned Pairs. arXiv:2408.09432.
- Yaniv, Z.; Lowekamp, B. C.; Johnson, H. J.; and Beare, R. 2018. SimpleITK image-analysis notebooks: a collaborative environment for education and reproducible research. *Journal of digital imaging*, 31(3): 290–303.
- Yoo, T. K.; Choi, J. Y.; and Kim, H. K. 2020. A generative adversarial network approach to predicting postoperative appearance after orbital decompression surgery for thyroid eye disease. *Computers in Biology and Medicine*, 118: 103628.
- Yoon, J. S.; Zhang, C.; Suk, H.-I.; Guo, J.; and Li, X. 2023. *SADM: Sequence-Aware Diffusion Model for Longitudinal Medical Image Generation*, 388–400. Springer Nature Switzerland. ISBN 9783031340482.
- Zhang, T.; Cheng, J.; Fu, H.; Gu, Z.; Xiao, Y.; Zhou, K.; Gao, S.; Zheng, R.; and Liu, J. 2019. Noise adaptation generative adversarial network for medical image analysis. *IEEE transactions on medical imaging*, 39(4): 1149–1159.
- Zhu, J.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. *CoRR*, abs/1703.10593.